

NOvA Operations - Bug # 1068: DCM Timing Chain Related Issue

Status:	New	Priority:	High
Author:	Peter Shanahan	Category:	
Created:	03/17/2011	Assigned to:	
Updated:	04/18/2011	Due date:	
component:	base		
Subject:	DCM Timing Chain Related Issue		
Description:	Generally, the symptom is several DCMs will go saturated in CPU use. They (almost?) without exception will be the last DCMs in the timing chain. Things to check: Does the first bad DCM have a timing input link problem?		

History

03/30/2011 02:04 pm - Leon Mualem

- File *Ganglia___DCM_Cluster_Report_2011_03_30.pdf* added

<http://dbweb2.fnal.gov:8080/ECL/nova/E/show?e=4795>

The above logbook entry reports some information on what happened last night that appears to be related to this problem.

I'll copy some of it in here for completeness.

Category: General

Posted: 03/30/2011 13:57:22

By: Leon Mualem (mualem)

Tags: Commissioning DCM

Text: There were some issues I noticed in the ganglia plots this morning that pertain to the overnight run. At about 4:20 AM there was some problem that disrupted the previously stable DAQ. Report that follows goes along the timing chain, which we suspect might be part of the problem:

DCM 6 and 8 were undisturbed, and appeared to continue to send data without trouble.

DCM-11 Stopped sending significant data, and CPU and running processes dropped almost to zero. There may have been some corrupt microslices from it, but the DCM is currently down, so it does not show this metric in ganglia.

DCM-12 stopped sending significant data, CPU saturated, load went way up to ~5, and reported corrupt microslices at ~125Hz.

DCM-09 stopped sending significant data, CPU and load dropped to essentially zero, reported a few corrupt microslices, then went quiet.

DCM-13 stopped sending significant data, CPU went very low, load stayed low, but 1min load went high, up to about 8 (gray far above blue line), generated some corrupt microslices initially, and then nothing.

So things along the timing chain all went bad past the first 2 DCMs, but they appear to go bad in different ways.

You can look it up in ganglia yourself, or see some plots in the attached pdf.

03/30/2011 03:17 pm - Leon Mualem

- File *Ganglia___DCM_Cluster_Report_2011_03_30_0800.pdf* added

Posted entry 4796, along these same lines.

Category: General

Posted: 03/30/2011 15:11:06

By: Leon Mualem (mualem)

Tags: Commissioning DCM

Relates to: 4795, Leon Mualem (mualem), 03/30/2011 13:57:22

[show thread](#)

Text: There were some issues I noticed in the ganglia plots this morning that pertain to running this morning. At about 9:00 AM there was some problem that disrupted the previously stable DAQ. Report that follows goes along the timing chain, which we suspect might be part of the problem:

DCM 6 appears undisturbed, and appeared to continue to send data without trouble.

DCM 8 data rate drops, CPU load goes saturated for about 10 minutes and then falls to essentially nothing. IT probably has some corrupt microslices, at the ~100hz level, but it is currently down, so I can't see it on ganglia.

DCM-11 Date rate dropped, but appeared to still be sending data at a reduced rate, ~500kB/s instead of 1.5MB/s it was previously sending. CPU and running processes jumped significantly. There may have been some corrupt microslices from it, but the DCM is currently down, so it does not show this metric in ganglia.

DCM-12 stopped sending significant data, CPU saturated, load went way up to ~5, and reported corrupt microslices at ~50Hz.

DCM-09 stopped sending significant data, CPU and load went up significantly, saturating CPU and load at ~3, reported corrupt microslices ~125Hz.

DCM-13 data rate went down significantly, CPU went saturated, load went up, generated ~100Hz corrupt microslices.

So things along the timing chain all went bad past the first DCM. In this case they seem to go bad mostly in the same way, except for a couple that were still sending some data.

You can look it up in ganglia yourself, or see some plots in the attached pdf.

03/31/2011 04:20 pm - Leon Mualem

Looking at the log on dcm-09 from this morning at the start of the corruption, I find a few strange things in the timing. Maybe it's just that I don't understand.

Early data show time boundaries that end in 3196. Later they would end in all 0s.

08:06:52.476841 2 Sending millislice with SEQ number 6348900, tstart 2513460000003196, size 2701, nmicro 100

08:06:52.954711 2 Sending millislice with SEQ number 6349000, tstart 2513460032003196, size 2917, nmicro 100

08:06:53.456193 2 Sending millislice with SEQ number 6349100, tstart 2513460064003196, size 2779, nmicro 100

08:06:53.955992 2 Sending millislice with SEQ number 6349200, tstart 2513460096003196, size 2680, nmicro 100

Then corruption started and the dcmapp would have to catch up, jumping 5 or 10 seconds each time.

When some real data shows up it has a different time. It seems like the time should always be adding an even multiple, 3200*100, so I don't understand why it is not on the proper boundary when it is skipping.

08:13:45.078295 2 Sending millislice with SEQ number 6430900, tstart 2513486240000000, size 19, nmicro 0

08:13:45.342750 2 Sending millislice with SEQ number 6431000, tstart 2513486272000000, size 19, nmicro 0

08:13:46.206829 2 Sending millislice with SEQ number 6431100, tstart 2513486304022396, size 13417, nmicro 2258

08:13:46.499846 2 Sending millislice with SEQ number 6431200, tstart 2513486336000000, size 19, nmicro 0

08:13:46.748158 2 Sending millislice with SEQ number 6431300, tstart 2513486368000000, size 19, nmicro 0

Other odd features:

Occasionally the problem with the corrupt DT sometimes does not have a multiple of 3200. It can be off by 4. It could be -4, 4, -12804, 3196, 86396,

..

03/31/2011 05:34 pm - Leon Mualem

And some more data from problems this afternoon 2011-03-31-1700hrs.

Corruption formed on DCM 12,9,13 on one instance, and then on all, starting from dcm-06.

04/15/2011 02:39 pm - Leon Mualem

Since I haven't updated this in a couple of weeks, and since we talked about things, and had a few email exchanges about this, I thought I'd capture a few of the ideas and exchanges here.

- 1) It seems that it is really a timing-related issue somehow, in that several of the DCMs jumped and had corruption at the same time.
- 2) The time jump can be forward or backward.
- 3) It seems like ALL FEBs jump at once, since the microslice rate does not change from the expected 20kHz.

And an email trail from Sue:

I've attached the time markers converted to hex in this e-mail. I didn't notice before, but all 3 dcms show exactly the same "Last" Time Marker before flipping over to corrupt time markers. Maybe that means something.

Regards,

-Sue

> > Here's a set of snips from this morning's time marker corruption on 3 dcms (dcm-03,dcm-08, and dcm-11). All 3

> > reported corruption at 7:15 AM and all 3 experienced a jump back on the timing marker.

> > <http://homepages.spa.umn.edu/~schubert/tmp/rick.log>

> > The times are in decimal in novatime - I'll convert them over to hex. -Sue

dcm-03:

07:15:00.357166 2 Sending millislice with SEQ number 5024500, tstart 2590675232003196, size 2902, nmicro 100 %MSG-w
DcmReaderDataCorrupt: dcm-3-01-01 14-Apr-2011 07:15:00 CDT MF-online DcmReader.cpp:969 0)

DcmReader::_checkMicroSliceCorruption - TM before Prev. Skip TM 2590674093778300. Last TM 2590675260870396. dT
-1167092096.

%MSG

Translates to in hex:

Last TM: 2590675260870396/0x9343496A70EFC

Curr TM: 2590674093778300/0x934345116A57C

dcm-08:

07:15:00.612924 2 Sending millislice with SEQ number 5024500, tstart 2590675232003196, size 2941, nmicro 100

%MSG-w DcmReaderDataCorrupt: dcm-3-01-02 14-Apr-2011 07:15:01 CDT MF-online DcmReader.cpp:969

0) DcmReader::_checkMicroSliceCorruption - TM before Prev. Skip TM 2590674093874300. Last TM 2590675260870396. dT
-1166996096.

%MSG

Last TM: 2590675260870396/0x9343496A70EFC

Curr TM: 2590674093874300/0x9343451181C7C

dcm-11:

07:15:00.505087 2 Sending millislice with SEQ number 5024500, tstart 2590675232003196, size 1660, nmicro 100

%MSG-w DcmReaderDataCorrupt: dcm-3-01-03 14-Apr-2011 07:15:00 CDT MF-online DcmReader.cpp:969

0) DcmReader::_checkMicroSliceCorruption - TM before Prev. Skip TM 2590674093797500. Last TM 2590675260870396. dT
-1167072896.

%MSG

Last TM: 2590675260870396/0x9343496A70EFC

Curr TM: 2590674093797500/0x934345116F07C

And the referred to log:

dcm-03:

07:14:59.363227 2 Sending millislice with SEQ number 5024300, tstart 2590675168003196, size 2818, nmicro 100

07:14:59.877401 2 Sending millislice with SEQ number 5024400, tstart 2590675200003196, size 2905, nmicro 100

07:15:00.357166 2 Sending millislice with SEQ number 5024500, tstart 2590675232003196, size 2902, nmicro 100

%MSG-w DcmReaderDataCorrupt: dcm-3-01-01 14-Apr-2011 07:15:00 CDT MF-online DcmReader.cpp:969

0) DcmReader::_checkMicroSliceCorruption - TM before Prev. Skip TM 2590674093778300. Last TM 2590675260870396. dT -1167092096.

%MSG

%MSG-w DcmReaderDataCorrupt: dcm-3-01-01 14-Apr-2011 07:15:00 CDT MF-online DcmReader.cpp:969

1) DcmReader::_checkMicroSliceCorruption - TM before Prev. Skip TM 2590674093775100. Last TM 2590675260870396. dT -1167095296.

%MSG

%MSG-w DcmReader: dcm-3-01-01 14-Apr-2011 07:32:30 CDT MF-online DcmReader.cpp:668

DcmReader::getData has fallen behind wall clock by 1.04806e+06 ms. Will jump ahead to catch up.

%MSG

%MSG-i DcmReader: dcm-3-01-01 14-Apr-2011 07:32:30 CDT MF-online DcmReader.cpp:701

DcmReader::getData - Jump by 1171 SEQ numbers to milliSliceSEQ 5025762, Time 2590675635840000.

%MSG

%MSG-w DcmReaderDataCorrupt: dcm-3-01-01 14-Apr-2011 07:32:30 CDT MF-online DcmReader.cpp:969

0) DcmReader::getData - new TM not the expected DT from previous. TM 2590675636123900. Last TM 2590675636101500. dT 22400.

%MSG

%MSG-w DcmReaderDataCorrupt: dcm-3-01-01 14-Apr-2011 07:32:30 CDT MF-online DcmReader.cpp:969

1) DcmReader::getData - new TM not the expected DT from previous. TM 2590675636107900. Last TM 2590675636123900. dT -16000.

%MSG

%MSG-w DcmReader: dcm-3-01-01 14-Apr-2011 07:32:30 CDT MF-online DcmReader.cpp:668

DcmReader::getData has fallen behind wall clock by 1.04248e+06 ms. Will jump ahead to catch up.

dcm-08:

07:15:00.113821 2 Sending millislice with SEQ number 5024400, tstart 2590675200003196, size 3247, nmicro 100

07:15:00.612924 2 Sending millislice with SEQ number 5024500, tstart 2590675232003196, size 2941, nmicro 100

%MSG-w DcmReaderDataCorrupt: dcm-3-01-02 14-Apr-2011 07:15:01 CDT MF-online DcmReader.cpp:969

0) DcmReader::_checkMicroSliceCorruption - TM before Prev. Skip TM 2590674093874300. Last TM 2590675260870396. dT -1166996096.

%MSG

%MSG-w DcmReaderDataCorrupt: dcm-3-01-02 14-Apr-2011 07:15:01 CDT MF-online DcmReader.cpp:969

1) DcmReader::_checkMicroSliceCorruption - TM before Prev. Skip TM 2590674093871100. Last TM 2590675260870396. dT -1166999296.

%MSG

07:15:01.337719 2 Sending millislice with SEQ number 5024600, tstart 2590675264000000, size 19, nmicro 0

07:15:02.842635 2 Sending millislice with SEQ number 5024700, tstart 2590675296000000, size 19, nmicro 0

07:15:04.485800 2 Sending millislice with SEQ number 5024800, tstart 2590675328000000, size 19, nmicro 0

07:15:06.778562 2 Sending millislice with SEQ number 5024900, tstart 2590675360000000, size 19, nmicro 0

07:15:09.109866 2 Sending millislice with SEQ number 5025000, tstart 2590675392000000, size 19, nmicro 0

%MSG-w DcmReader: dcm-3-01-02 14-Apr-2011 07:15:10 CDT MF-online DcmReader.cpp:668

DcmReader::getData has fallen behind wall clock by 5065 ms. Will jump ahead to catch up.

%MSG

%MSG-i DcmReader: dcm-3-01-02 14-Apr-2011 07:15:10 CDT MF-online DcmReader.cpp:701

DcmReader::getData - Jump by 9329 SEQ numbers to milliSliceSEQ 5034366, Time 2590678389120000.

%MSG

%MSG-w DcmReaderDataCorrupt: dcm-3-01-02 14-Apr-2011 07:15:10 CDT MF-online DcmReader.cpp:969

0) DcmReader::_checkMicroSliceCorruption - TM before Prev. Skip TM 2590674094181500. Last TM 2590678389199996. dT -4295018496.

%MSG

%MSG-w DcmReaderDataCorrupt: dcm-3-01-02 14-Apr-2011 07:15:10 CDT MF-online DcmReader.cpp:969

1) DcmReader::_checkMicroSliceCorruption - TM before Prev. Skip TM 2590674094207100. Last TM 2590678389199996. dT -4294992896.

%MSG

%MSG-w DcmReader: dcm-3-01-02 14-Apr-2011 07:15:56 CDT MF-online DcmReader.cpp:668

dcm-11:

07:14:58.997587 2 Sending millislice with SEQ number 5024200, tstart 2590675136003196, size 1417, nmicro 100

07:14:59.497236 2 Sending millislice with SEQ number 5024300, tstart 2590675168003196, size 1384, nmicro 100

07:15:00.004060 2 Sending millislice with SEQ number 5024400, tstart 2590675200003196, size 1381, nmicro 100

07:15:00.505087 2 Sending millislice with SEQ number 5024500, tstart 2590675232003196, size 1660, nmicro 100

%MSG-w DcmReaderDataCorrupt: dcm-3-01-03 14-Apr-2011 07:15:00 CDT MF-online DcmReader.cpp:969

0) DcmReader::_checkMicroSliceCorruption - TM before Prev. Skip TM 2590674093797500. Last TM 2590675260870396. dT -1167072896.

%MSG

%MSG-w DcmReaderDataCorrupt: dcm-3-01-03 14-Apr-2011 07:15:00 CDT MF-online DcmReader.cpp:969

1) DcmReader::_checkMicroSliceCorruption - TM before Prev. Skip TM 2590674093848700. Last TM 2590675260870396. dT -1167021696.

%MSG

%MSG-w DcmReader: dcm-3-01-03 14-Apr-2011 07:17:53 CDT MF-online DcmReader.cpp:668

DcmReader::getData has fallen behind wall clock by 170325 ms. Will jump ahead to catch up.

%MSG

%MSG-i DcmReader: dcm-3-01-03 14-Apr-2011 07:17:53 CDT MF-online DcmReader.cpp:701

DcmReader::getData - Jump by 1025 SEQ numbers to milliSliceSEQ 5025616, Time 2590675589120000.

%MSG

%MSG-w DcmReaderDataCorrupt: dcm-3-01-03 14-Apr-2011 07:17:53 CDT MF-online DcmReader.cpp:969

0) DcmReader::getData - new TM not the expected DT from previous. TM 2590675589339900. Last TM 2590675589231100. dT 108800.

%MSG

%MSG-w DcmReaderDataCorrupt: dcm-3-01-03 14-Apr-2011 07:17:53 CDT MF-online DcmReader.cpp:969

04/18/2011 02:17 pm - Leon Mualem

Recently with all the scintillator filling and top-up operations going on it appears that when the pumps start or turn off there is a strong correlation with these corruption events. My suspicion is that this would be some common mode offset, or some noise associated with start-up that is causing glitching on the timing system lines. It might be that this is showing up since all the DCMs are running off the floating 24V supply, but the TDUs are running off a line-power supply. I believe some thoughts about this being a potential problem are already being ruminated upon. My pet theory is that we get an extra "sync" pulse that glitches the clocks into the previously set state. I believe it shouldn't happen, but maybe it is. It may be that some time at the detector experimenting with turning things on and off and watching signals would be in order.

Files

Ganglia__DCM_Cluster_Report_2011_03_30.pdf	282.9 KB	03/30/2011	Leon Mualem
Ganglia__DCM_Cluster_Report_2011_03_30_0800.pdf	295.8 KB	03/30/2011	Leon Mualem